

	Type	L #	Hits	Search Text	DBs	Time Stamp
1	BRS	L1	0	thread and ins truction adj stream	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/04/17 14:05
2	BRS	L2	1099	thread and instruction adj stream	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/04/17 14:05
3	BRS	L3	450	thread same instruction adj stream	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/04/17 14:05
4	BRS	L4	0	3 AND register sane data adj element	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/04/17 14:06
5	BRS	L5	18	3 AND register same data adj element	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/04/17 14:11
6	BRS	L6	406	3 AND register	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/04/17 14:11

7	BRS	L7	388	6 not 5	US- PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/04/17 14:11
---	-----	----	-----	---------	--	---------------------

	Type	L #	Hits	Search Text	DBs	Time Stamp
8	BRS	L8	244	7 and register adj file	US- PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/04/17 15:06
9	BRS	L9	13045639	@ad<"19950816"	US- PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/04/17 14:12
10	BRS	L10	11	7 and 9	US- PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/04/17 14:12
11	BRS	L11	49	7 and register adj file and VLIW	US- PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/04/17 15:10
12	BRS	L12	2	11 and 9	US- PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/04/17 15:10
13	BRS	L13	244	7 and register adj file	US- PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/04/17 15:10

14	BRS	L14	6	13 and 9	US- PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/04/17 15:10
----	-----	-----	---	----------	--	---------------------

	Type	L #	Hits	Search Text	DBs	Time Stamp
15	BRS	L15	4	14 not 12	US- PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	2006/04/17 15:10



USPTO

[Subscribe](#) (Full Service) [Register](#) (Limited Service, Free) [Login](#)

Search: ☒ The ACM Digital Library ☐ The Guide

"data Path" and "register file" and "data element" and therad and

SEARCH

THE ACM DIGITAL LIBRARY



[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used [data Path](#) and [register file](#) and [data element](#) and [therad](#) and [round robin](#) and [pipeline](#) and [external interface](#) and [instruction stream](#)

Found 630 of 175,083

Sort results by

relevance

[Save results to a Binder](#)

[Try an Advanced Search](#)

Display results

expanded form

[Search Tips](#)

[Try this search in The ACM Guide](#)

☐ Open results in a new window

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale ☐ ☐ ☐ ☐ ☐

1 [A performance analysis of PIM, stream processing, and tiled processing on memory-intensive signal processing kernels](#)



Jinwoo Suh, Eun-Gyu Kim, Stephen P. Crago, Lakshmi Srinivasan, Matthew C. French  
May 2003 **ACM SIGARCH Computer Architecture News , Proceedings of the 30th annual international symposium on Computer architecture ISCA '03**, Volume 31 Issue 2

**Publisher:** ACM Press

Full text available: pdf(239.50 KB) Additional Information: [full citation](#), [abstract](#), [references](#)

Trends in microprocessors of increasing die size and clock speed and decreasing feature sizes have fueled rapidly increasing performance. However, the limited improvements in DRAM latency and bandwidth and diminishing returns of increasing superscalar ILP and cache sizes have led to the proposal of new microprocessor architectures that implement processor-in- memory, stream processing, and tiled processing. Each architecture is typically evaluated separately and compared to a baseline architectu ...

2 [PSCP: a scalable parallel ASIP architecture for reactive systems](#)

A. Pyttel, A. Sedlmeier, C. Veith  
February 1998 **Proceedings of the conference on Design, automation and test in Europe**

**Publisher:** IEEE Computer Society

Full text available: pdf(208.55 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citing](#), [index](#)



[Publisher Site](#)

[terms](#)

We describe a Codesign approach based on a parallel and scalable ASIP architecture, which is suitable for the implementation of reactive systems. The specification language of our approach is extended statecharts. Our ASIP architecture is scalable with respect to the number of processing elements as well as parameters such as bus widths and register file sizes. Instruction sets are generated from a library of components covering a spectrum of space/time trade-off alternatives. Our approach featu ...

**Keywords:** FPGA, application-specific, statechart, modular


3 [Special session on reconfigurable computing: The happy marriage of architecture and application in next-generation reconfigurable systems](#)

 Ingrid Verbaauwhede, Patrick Schaumont  
April 2004 **Proceedings of the 1st conference on Computing frontiers**

**Publisher:** ACM Press  
Full text available:  [pdf\(398.28 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)


New applications and standards are first conceived only for functional correctness and without concerns for the target architecture. The next challenge is to map them onto an architecture. Embedding such applications in a portable, low-energy context is the art of molding it onto an energy-efficient target architecture combined with an energy efficient execution. With a reconfigurable architecture, this task becomes a two-way process where the architecture adapts to the application and vice-vers ...

**Keywords:** embedded, real-time systems

4 Space-time scheduling of instruction-level parallelism on a raw machine   
Walter Lee, Rajeev Barua, Matthew Frank, Devabhaktuni Srikrishna, Jonathan Babb, Vivek Sarkar, Saman Amarasinghe  
October 1998 **ACM SIGPLAN Notices , ACM SIGOPS Operating Systems Review , Proceedings of the eighth international conference on Architectural support for programming languages and operating systems ASPLOS-VIII**, Volume 33 , 32 Issue 11 , 5


**Publisher:** ACM Press  
Full text available:  [pdf\(1.79 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Increasing demand for both greater parallelism and faster clocks dictate that future generation architectures will need to decentralize their resources and eliminate primitives that require single cycle global communication. A Raw microprocessor distributes all of its resources, including instruction streams, register files, memory ports, and ALUs, over a pipelined two-dimensional mesh interconnect, and exposes them fully to the compiler. Because communication in Raw machines is distributed, com ...

5 Exploiting choice: instruction fetch and issue on an implementable simultaneous multithreading processor   
Dean M. Tullsen, Susan J. Eggers, Joel S. Emer, Henry M. Levy, Jack L. Lo, Rebecca L. Stamm  
May 1996 **ACM SIGARCH Computer Architecture News , Proceedings of the 23rd annual international symposium on Computer architecture ISCA '96**, Volume 24 Issue 2

**Publisher:** ACM Press  
Full text available:  [pdf\(1.48 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Simultaneous multithreading is a technique that permits multiple independent threads to issue multiple instructions each cycle. In previous work we demonstrated the performance potential of simultaneous multithreading, based on a somewhat idealized model. In this paper we show that the throughput gains from simultaneous multithreading can be achieved *without* extensive changes to a conventional wide-issue superscalar, either in hardware structures or sizes. We present an architecture for s ...

6 Strategies for achieving improved processor throughput   
Matthew K. Farrens, Andrew R. Pleszkun  
April 1991 **ACM SIGARCH Computer Architecture News , Proceedings of the 18th annual international symposium on Computer architecture ISCA '91**, Volume 19 Issue 3  
**Publisher:** ACM Press

Full text available:  [pdf\(742.44 KB\)](#) Additional Information: [full citation](#), [references](#), [citing](#)s, [index terms](#)

7 Pipeline Architecture



C. V. Ramamoorthy, H. F. Li  
March 1977 **ACM Computing Surveys (CSUR)**, Volume 9 Issue 1

**Publisher:** ACM Press


Full text available:  [pdf\(3.53 MB\)](#) Additional Information: [full citation](#), [references](#), [citing](#)s, [index terms](#)

8 Polygon rendering on a stream architecture



John D. Owens, William J. Dally, Ujval J. Kapasi, Scott Rixner, Peter Mattson, Ben Mowery  
August 2000 **Proceedings of the ACM SIGGRAPH/EUROGRAPHICS workshop on Graphics hardware**

**Publisher:** ACM Press

Full text available:  [pdf\(161.65 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citing](#)s, [index terms](#)

The use of a programmable stream architecture in polygon rendering provides a powerful mechanism to address the high performance needs of today's complex scenes as well as the need for flexibility and programmability in the polygon rendering pipeline. We describe how a polygon rendering pipeline maps into data streams and kernels that operate on streams, and how this mapping is used to implement the polygon rendering pipeline on Imagine, a programmable stream processor. We compare our resul ...


**Keywords:** OpenGL, SIMD, graphics hardware, kernels, media processors, polygon rendering, stream architecture, stream processing, streams

9 A survey of processors with explicit multithreading



Theo Ungerer, Borut Robič, Jurij Silc  
March 2003 **ACM Computing Surveys (CSUR)**, Volume 35 Issue 1

**Publisher:** ACM Press

Full text available:  [pdf\(920.16 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citing](#)s, [index terms](#)

Hardware multithreading is becoming a generally applied technique in the next generation of microprocessors. Several multithreaded processors are announced by industry or already into production in the areas of high-performance microprocessors, media, and network processors. A multithreaded processor is able to pursue two or more threads of control in parallel within the processor pipeline. The contexts of two or more threads of control are often stored in separate on-chip register sets. Unused i ...

**Keywords:** Blocked multithreading, interleaved multithreading, simultaneous multithreading

10 On pipelining dynamic instruction scheduling logic



Jared Stark, Mary D. Brown, Yale N. Patt  
December 2000 **Proceedings of the 33rd annual ACM/IEEE international symposium on Microarchitecture**

**Publisher:** ACM Press

Full text available:  [pdf\(128.82 KB\)](#)  [ps\(543.84 KB\)](#) Additional Information: [full citation](#), [references](#), [citing](#)s, [index terms](#)





- 11

GPGPU: general purpose computation on graphics hardware

David Luebke, Mark Harris, Jens Krüger, Tim Purcell, Naga Govindaraju, Ian Buck, Cliff Woolley, Aaron Lefohn

August 2004 **Proceedings of the conference on SIGGRAPH 2004 course notes GRAPH '04**

**Publisher:** ACM Press

Full text available: [pdf\(63.03 MB\)](#) Additional Information: [full citation](#), [abstract](#)

The graphics processor (GPU) on today's commodity video cards has evolved into an extremely powerful and flexible processor. The latest graphics architectures provide tremendous memory bandwidth and computational horsepower, with fully programmable vertex and pixel processing units that support vector operations up to full IEEE floating point precision. High level languages have emerged for graphics hardware, making this computational power accessible. Architecturally, GPUs are highly parallel s ...

☐
- 12

Session 17: architecture: Sunder: a programmable hardware prefetch architecture for numerical loops

Tzi-cker Chiueh

November 1994 **Proceedings of the 1994 ACM/IEEE conference on Supercomputing**

**Publisher:** ACM Press

Full text available: [pdf\(922.38 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

*Beyond data caching, data prefetching is by far the most effective way to address the memory access bottleneck associated with high-performance processors. This is particularly true for scientific programs whose working sets cannot be easily fit into the on-chip data cache. This paper proposes a new data prefetching architecture called **Sunder**, which combines the flexibility and accurateness of software prefetching and the transparency and low-overhead of hardware prefetching. Th ...*

☐

13

Dataflow Mini-Graphs: Amplifying Superscalar Capacity and Bandwidth

Anne Bracy, Prashant Prahlaad, Amir Roth

December 2004 **Proceedings of the 37th annual IEEE/ACM International Symposium on Microarchitecture MICRO 37**

**Publisher:** IEEE Computer Society

Full text available: [pdf\(189.31 KB\)](#) Additional Information: [full citation](#), [abstract](#), [citations](#)

A mini-graph is a dataflow graph that has an arbitrary internal size and shape but the interface of a singleton instruction: two register inputs, one register output, a maximum of one memory operation, and a maximum of one (terminal) control transfer. Previous work has exploited dataflow sub-graphs whose execution latency can be reduced via programmable FPGA-style hardware. In this paper we show that mini-graphs can improve performance by amplifying the bandwidths of a superscalar processor's st ...

☐

14

Architectures: A programmable vertex shader with fixed-point SIMD datapath for low power wireless applications

Ju-Ho Sohn, Ramchan Woo, Hoi-Jun Yoo

August 2004 **Proceedings of the ACM SIGGRAPH/EUROGRAPHICS conference on Graphics hardware**

**Publisher:** ACM Press

Full text available: [pdf\(427.49 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The real time 3D graphics becomes one of the attractive applications for 3G wireless terminals although their battery lifetime and memory bandwidth limit the system

☐

http://portal.acm.org/results.cfm?coll=ACM&dl=ACM&CFID=69633554&CFTOKEN=2364903

4/17/06

resources for graphics processing. Instead of using the dedicated hardware engine with complex functions, we propose an efficient hardware architecture of low power vertex shader with programmability. Our architecture includes the following three features: 1) a fixed-point SIMD datapath to exploit parallelism in vertex process ...

15 The M-Machine multicomputer



Marco Fillo, Stephen W. Keckler, William J. Dally, Nicholas P. Carter, Andrew Chang, Yevgeny Gurevich, Whay S. Lee

December 1995 **Proceedings of the 28th annual international symposium on Microarchitecture**

**Publisher:** IEEE Computer Society Press

Full text available: [pdf\(1.29 MB\)](#) Additional Information: [full citation](#), [references](#), [citing](#), [index terms](#)

16 Graphics rendering architecture for a high performance desktop workstation



Chandlee B. Harrell, Farhad Fouladi

September 1993 **Proceedings of the 20th annual conference on Computer graphics and interactive techniques**

**Publisher:** ACM Press

Full text available: [pdf\(346.15 KB\)](#) Additional Information: [full citation](#), [references](#), [citing](#), [index terms](#)

17 Compilation: Cluster assignment of global values for clustered VLIW processors



Andrei Terechko, Erwan Le Thénaff, Henk Corporaal

October 2003 **Proceedings of the 2003 international conference on Compilers, architecture and synthesis for embedded systems**

**Publisher:** ACM Press

Full text available: [pdf\(330.94 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In this paper high-level language (HLL) variables that are alive in a whole HLL function, across multiple scheduling units, are termed as global values. Due to their long live ranges and, hence, large impact on the schedule, the global values require different compiler optimizations than local values, which span across only one scheduling unit. The instruction scheduler for a clustered ILP processor, which is responsible for cluster assignment of operations and variables, faces a difficult probl ...

**Keywords:** ILP, VLIW, cluster assignment, compiler, instruction scheduler, register allocation

18 Ray tracing vs. scan conversion: Comparing Reyes and OpenGL on a stream architecture



John D. Owens, Brucek Khailany, Brian Towles, William J. Dally

September 2002 **Proceedings of the ACM SIGGRAPH/EUROGRAPHICS conference on Graphics hardware**


**Publisher:** Eurographics Association

Full text available: [pdf\(136.72 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citing](#), [index terms](#)

The OpenGL and Reyes rendering pipelines each render complex scenes from similar scene descriptions but differ in their internal pipeline organizations. While the OpenGL organization has dominated hardware architectures over the past twenty years, a Reyes organization differs in several important ways from OpenGL, including a shader coordinate system that supports coherent texture accesses, a single shader in the vertex stage, and tessellation and sampling instead of triangle rasterization. Hardw ...

19

A low-power memory hierarchy for a fully programmable baseband processor

 Wolfgang Raab, Hans-Martin Bluethgen, Ulrich Ramacher  
June 2004 **Proceedings of the 3rd workshop on Memory performance issues: in conjunction with the 31st international symposium on computer architecture WPMP '04**

**Publisher:** ACM Press  
Full text available:  [pdf\(431.55 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)


Future terminals for wireless communication not only must support multiple standards but execute several of them concurrently. To meet these requirements, flexibility and ease of programming of integrated circuits for digital baseband processing are increasingly important criteria for the deployment of such devices, while power consumption and area of the devices remain as critical as in the past. The paper presents the architecture of a fully programmable system-on-chip for digital signal proces ...

**Keywords:** baseband processor, low-power memory, memory hierarchy, multi-tasked processor, task interleaving

20

Trident: a scalable architecture for scalar, vector, and matrix operations

Mostafa I. Soliman, Stanislav G. Sedukhin  
January 2002 **Australian Computer Science Communications , Proceedings of the seventh Asia-Pacific conference on Computer systems architecture - Volume 6 CRPITS '02**, Volume 24 Issue 3

**Publisher:** Australian Computer Society, Inc. , IEEE Computer Society Press  
Full text available:  [pdf\(814.51 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Within a few years it will be possible to integrate a billion transistors on a single chip. At this integration level, we propose using a high level ISA to express parallelism to hardware instead of using a huge transistor budget to dynamically extract it. Since the fundamental data structures for a wide variety of applications are scalar, vector, and matrix, our proposed Trident processor extends the classical vector ISA with matrix operations. The Trident processor consists of a set of paralle ...

**Keywords:** data parallelism, parallel processing, ring register file, scalable hardware, vector/matrix processing